
TRANSCRIPTS

Bubble, Rubble, Finance in Trouble?

Andrew W. Lo

In this talk, I review the implications of the recent rise and fall of the technology sector for traditional financial theories and their behavioral alternatives. Although critics of the Efficient Markets Hypothesis argue that markets are driven by fear and greed, not fundamentals, recent research in the cognitive neurosciences suggest that these two perspectives are opposite sides of the same coin. I propose a new paradigm for financial economics that focuses more on the evolutionary biology and ecology of markets rather than the more traditional physicists' view. By marrying the principles of evolution to Herbert Simon's notion of "satisficing," I argue that much of what behavioralists cite as counter-examples to economic rationality—loss aversion, overconfidence, overreaction, mental accounting, and other behavioral biases—are, in fact, consistent with an evolutionary model of rational agents learning to adapt to their environment via satisficing heuristics.

I'd like to begin by thanking David Dreman and Arnie Wood for inviting me to speak at this fascinating conference. Although I'm not a psychologist or behavioral finance expert by training, much of my current research interests lies in this area now, and I'll be discussing some of those interests over the next 45 minutes or so. This luncheon talk provides me with a wonderful opportunity to speak somewhat less formally and a bit more expansively about several issues surrounding behavioral finance and how they relate to the apparent bubble in technology stocks that we have experienced over the past few years. Given the occasion, my comments will be less than rigorous, but I've often been reminded by my colleagues that along with rigor usually goes mortis, so I trust no one will object in this case!

What I thought I would do instead is to describe part of the personal journey that I've taken in resolving the controversy surrounding behavioral finance and in attempting to reconcile traditional financial asset-pricing models with what's been happening in the stock mar-

ket over the past several years. As part of that journey, I think I've been able to address some of the challenges that David Dreman raised in his opening remarks yesterday, challenges that, at one point, I found enormously difficult to make sense of in the context of modern financial economics. It's only recently that I think I've finally come to a tentative reconciliation, at least in my own mind, of these two apparently disparate lines of inquiry.

The Three P's of Investment Management

Let me start by answering the title of my talk, "Bubble, Rubble, Finance in Trouble?" and state at the outset—with apologies to William Shakespeare—that despite the rubble left by the bursting of the technology bubble, modern finance is not in trouble. To see why, let me first provide a brief overview of what we consider to be the traditional or rational financial paradigm, which is summarized by what I call the "Three P's of Investment Management": prices, probabilities, and preferences. The three P's have their origins in one of the most basic and central ideas of modern economics, the principle of supply and demand. This principle states that the price of any commodity and the quantity traded are determined by the intersection of supply and demand curves, where the demand curve represents the schedule of quantities desired by consumers at various prices and the supply curve represents the schedule of quantities producers are willing to supply at various prices. The intersection of these two curves determines an "equilibrium," a price-quantity pair that satisfies both consumers and producers simultaneously. Any other price-quantity pair may serve one group's interests, but not the other's.

Even in this simple description of a market, all the elements of modern finance are present. The demand curve is the aggregation of many individual consumers' desires, each derived from optimizing an individual's preferences subject to a budget constraint that depends on prices and other factors (e.g., income, savings requirements, and borrowing costs). Similarly, the sup-

Andrew W. Lo is the Harris & Harris Group Professor and Director of the Laboratory for Financial Engineering at the Massachusetts Institute of Technology's Sloan School of Management.

Requests for reprints should be sent to: Andrew W. Lo, MIT Sloan School of Management, 50 Memorial Drive, E52-432, Cambridge, MA 02142. Email: alo@mit.edu

ply curve is the aggregation of many individual producers' outputs, each derived from optimizing an entrepreneur's preferences subject to a resource constraint that also depends on prices and other factors (e.g., costs of materials, wages, and trade credit). And probabilities affect both consumers and producers as they formulate their consumption and production plans through time and in the face of uncertainty—uncertain income, uncertain costs, and uncertain business conditions.

It's the interaction between prices, preferences, and probabilities that gives finance its richness and depth. Formal models of financial asset prices such as those of Merton (1973), Lucas (1978), and Breeden (1979), show precisely how the three P's simultaneously determine a "general equilibrium" in which demand equals supply across *all* markets in an uncertain world where individuals and corporations act rationally to optimize their own welfare. The three P's enter into any economic decision surrounding securities, or, for that matter, any other type of economic commodity. This naturally raises several questions about the three P's.

First, what exactly are preferences and how are they modeled? Financial economists do have answers to this question, but those answers are not completely satisfying—for example, utility functions and expected utility theory—despite the fact that they do yield interesting implications, many of which can be tested empirically. Unfortunately, most of the empirical tests seem to suggest important departures from these models of preferences.

Second, do people really behave this way, and are the three P's a *positive* or *normative* model of economic behavior? If individuals don't behave this way, *should* they? This is a much more open-ended and controversial topic. The finance literature has not really grappled with it head on just yet, but this is becoming more of a focus in some of the recent literature.

And then the last set of issues are those that I've been working on in my current research—issues that haven't yet made it into the mainstream of any literature—which have to do with a *theory* of behavioral finance. How do preferences arise? Where do preferences come from? And how do these preferences interact with market structure to give rise to the kinds of behavior that we observe?

I'll give concrete examples of all of these issues in just a few minutes, but I want begin by outlining what I see as the basic challenges to so-called "traditional" or "rational" finance. These questions really are the ones that should have been asked 20 or 30 years ago when some of the paradigms that we now take as canon in modern finance were first proposed. But, of course, much of the pathbreaking research that we now consider to be traditional finance only came about by considering the implications of those early paradigms, as opposed to struggling with some of the bigger-picture questions that this conference will be focusing on.

Estimating Probabilities

So let me turn to one outstanding issue. I think we all know that preferences are a problem, and that there are a variety of alternatives out there that try to capture departures from expected utility. But let me give you another example that has to do, instead, with probabilities. We heard a bit about probability matching from Mike Gazzaniga's wonderful talk last evening, and I want to come back to that in just a few minutes. But before we do, let me try to illustrate to you the challenges that humans have in processing probabilities. This is a very simple example that I'm sure many of you have seen, but if you haven't, I think you'll be surprised. The example has to do a hypothetical blood test for AIDS. This is a rather somber subject, but one that illustrates very clearly the point that I want to make about probabilities. Consider a blood test for AIDS that is 99 percent accurate. What I mean by "99 percent accurate" is that if you have AIDS and you take this blood test, the test will yield a positive result with 99 percent probability, and with 1 percent probability it will yield a negative result. And if you don't have AIDS, there is a 99 percent probability that the test will come back negative, and a 1 percent probability that it will yield a positive result. So, in other words, regardless of whether you have it or not, there is only a 1 percent chance that this test will make a mistake, either way. That's the sense in which this is a 99 percent accurate test.

The question that I want to ask you to consider is this: Suppose you take a blood test and, God forbid, it comes back positive. What's the probability that you have AIDS? Some of you have seen this example before, apparently, but if you haven't, just ask yourself what your reaction would be? What would your own assessment of your likelihood of having AIDS be? And how would you feel about such an outcome?

Now those of you with some probability training know that the way to calculate the answer is to use Bayes' rule. Specifically, what I've asked you to evaluate is the probability that you have AIDS given that you tested positive. But when you put this question to most people, even sophisticated individu-

als such as financial analysts, their response is typically "I have a 99 percent probability of having AIDS because, after all, the test is 99 percent accurate and I tested positive." Because the test is only incorrect with 1 percent probability, the incorrect conclusion that most people come to is that a positive test implies a 99 percent probability that the subject has AIDS.

The source of the confusion lies in the fact that the question I asked is *not* "what's the probability that you test positive given that you have AIDS," which is indeed 99 percent, but rather "what's the probability you have AIDS given that the blood test is positive." This

latter probability is related to the former by Bayes' rule in the following way:

$$\text{Prob}(\text{AIDS}|+) = \text{Prob}(+|\text{AIDS}) \times \frac{\text{Prob}(\text{AIDS})}{\text{Prob}(+)} \quad (1)$$

To evaluate this quantity, we need to plug in values for the three terms in the right side of (1). We know that $\text{Prob}(+|\text{AIDS})$ is 0.99 by assumption. $\text{Prob}(+)$ can be rewritten as:

$$\begin{aligned} \text{Prob}(+) &= \text{Prob}(+|\text{AIDS}) \times \text{Prob}(\text{AIDS}) \\ &+ \text{Prob}(+|\text{No AIDS}) \times [1 - \text{Prob}(\text{AIDS})]. \end{aligned} \quad (2)$$

Now if we're willing to make some rough approximations with respect to the unconditional probability of having AIDS—for example, there are approximately 250,000 individuals in the US with the AIDS virus out of a population of about 250 million—we have:

$$\text{Prob}(\text{AIDS}) \approx \frac{250,000}{250,000,000} = 0.1\% \quad (3)$$

This implies that:

$$\begin{aligned} \text{Prob}(\text{AIDS}|+) &= \text{Prob}(+|\text{AIDS}) \times \frac{\text{Prob}(\text{AIDS})}{\text{Prob}(+)} \\ &= 99\% \times \frac{0.1\%}{1.098\%} = 9.02\%. \end{aligned} \quad (3a)$$

It turns out that the probability that you have AIDS given a positive test is about 9 percent! That's 9 percent, not 99 percent! Isn't this a huge difference?

Now, of course, a 9 percent probability is nothing to celebrate because before you took the test, your base rate or unconditional probability of having AIDS was something like one-tenth of 1 percent, and now, having taken the blood test and testing positive, you've increased your probability almost by two orders of magnitude. But the point of this example is that even some of the most sophisticated investors get this question wrong because humans are simply not used to thinking in probabilistic terms. Probability is not really part of our prewiring, and as a result, we have to *learn* how to do these kinds of calculations.

Now, this gets to the point that was raised by David Dreman yesterday in his talk regarding analysts attempting to predict earnings. The prediction of earnings is, in my opinion, enormously more complex than dealing with a simple probability question such as the AIDS test example. And so are we really surprised that the forecast errors are so large? In some contexts, it's simply the nature of the beast. In fact, I'll go so far as to argue that there are good economic reasons why earnings forecasts *have* to be bad. Suppose that earning

forecasts were 99 percent accurate and further suppose that this implied that prices based on such forecasts were 99 percent accurate. That would yield huge profit opportunities that would be quickly bid away by many of you, after which the predictability would fall once again. So there are some very good economic reasons, institutional and structural reasons, that the forecast errors of financial forecasts are large. It's the nature of the beast, and it's a beautiful concept, but nevertheless it is one that we have to live with.

And by the way, we shouldn't apologize for this fact. A comedian said recently that he wanted to have a job as a TV weatherman because he couldn't think of any other job where you could be wrong so often and still show up to work and get paid! The same can be said for financial analysts, I think. A few years ago I started a small research project which I never published because I wasn't able to get enough data, but let me tell you about it anyway. I obtained some data about weather patterns from the National Weather Service and compared that to weather forecasts in the Boston area which I hand-collected for a few weeks, and believe me, the error rates of local forecasters for five-day forecasts were at least as bad as those of the typical financial analyst! But nobody really complains about it because we all know that's the nature of the weather—it's very hard to predict.

Now this doesn't excuse financial analysts from succumbing to the lure of the technology bubble, and I'll return to this issue in a few minutes, but I do want to emphasize that what we expect of financial analysts, and what we expect of many of you in the audience, is a very, very difficult and challenging set of tasks.

The Limits of Irrationality

I also mentioned that there are certain institutional limitations to these kinds of phenomena and restrictions on the extent to which behavioral biases can persist. Let me give you just one example that I think is well-known in academic circles but may not be as familiar to those of you in industry, and this is the so-called "Dutch book theorem." The idea is pretty simple—in certain circumstances, if people behave in ways that are inconsistent with the mathematics of probability theory, then there's money to be made! Here's a concrete illustration of the Dutch book theorem: Consider an event A which is defined as "the S&P 500 index drops by 2 percent next Monday." Suppose you ask an individual what his subjective probability is for such an event, and suppose the individual believes that A will occur with probability 50 percent, but at the same time, also believes that the probability A will *not* occur is 75 percent. This is clearly a violation of the basic axioms of probability theory—in particular, the probabilities of two mutually exclusive and exhaustive

events must sum to one—but there’s nothing that says that humans can’t violate the axioms of probability theory. In fact, quite often they do. What are the implications?

These subjective probabilities mean that the individual would be willing to take the following two bets. He would be willing to take a bet of getting paid a dollar if A occurs and paying a dollar if A doesn’t occur. That’s what it means to say that A is a 50/50 event. At the same time, if the individual thinks that the probability A doesn’t occur is 75 percent, he would be willing to take a bet where he gets paid a dollar if A doesn’t occur and pays three dollars if A does occur (75 percent probability of not-A is equivalent to 3/1 odds). In other words, the individual would be willing to take either side of both of the following bets:

$$B_1 : \begin{cases} \$1 \text{ if } A \\ -\$1 \text{ if not } A \end{cases} \quad (4a)$$

$$B_2 : \begin{cases} -\$3 \text{ if } A \\ \$1 \text{ if not } A \end{cases} \quad (4b)$$

Well, if an individual is truly willing to take either side of both of these bets, then that’s a happy situation for me because I’ll simply put \$50 on B_1 and \$25 on B_2 , and let’s see what happens. If A occurs, I lose \$50 on B_1 but I gain \$75 on B_2 , yielding a profit of \$25. If A doesn’t occur, I gain \$50 on B_1 and lose \$25 on B_2 (recall that B_2 is 3:1 odds), also yielding a profit of \$25. Regardless of the outcome, I’ve secured a profit of \$25, which is known as an “arbitrage” or a “money pump” in financial jargon. Obviously, such a happy circumstance is not sustainable and market forces—namely, arbitrageurs—will quickly react to take advantage of these discrepancies in probabilities until they force the odds to be in line with the axioms of probability theory. Therefore, there are limits to the degree and persistence of such irrational probability beliefs, and substantial incentives exist for those who can identify and exploit their occurrences. And if you don’t believe that you can have individuals with such inconsistent beliefs, an equivalent situation is where one group of individuals has one set of probability beliefs, another group has a conflicting set of beliefs, and under certain conditions, you can construct a very clever series of complex financial transactions to create these kind of arbitrages. This is what many so-called “relative value” hedge funds do as part of their daily operations, and quite a few have been handsomely rewarded for their efforts.

The point of this example is that while there are all sorts of behavioral regularities that are out there, and violations of what we consider to be rationality, at the same time, in certain cases, there are market forces that will bring prices back to their rational levels once

again. Such forces will tend to prevent these kinds of behavioral biases from becoming too significant. The bottom line question—the question that I’ve been struggling with for some time now—is how to reconcile the two? How can we reconcile the existence of behavioral biases with the existence of institutional and market forces that impose limits to irrationality, and how do they balance out?

In other words, how strong are such institutional and market forces? A good illustration of this conundrum is the collapse of the hedge fund Long-Term Capital Management (LTCM) in the summer of 1998. Some observers have concluded that the LTCM partners were “gunslingers,” irrationally risk-seeking, but in retrospect and upon closer examination, most researchers and industry professionals have come to the conclusion that they were acting quite rationally indeed. In fact, had they had more capital (recall that LTCM returned several billion dollars of capital to their investors in 1997), they would have been able to withstand the pressures of the market during the summer of 1998, and would have earned enormous sums of money from the large and, some would say, irrational swings in credit spreads in September and October of 1998. The question is which is stronger, the behavioral biases that we can see that can lead to this kind of a bubble or the market forces that will keep that tendency in check?

Psychology Versus Economics

And this brings me to my next point—the differences between psychology and economics. I want to spend a bit of time talking about that because my background and training is in economics. Now Mike Gazzaniga told us one amusing anecdote about economists yesterday; in fact, there are many jokes about economists. And the fact is that most of these jokes have an element of truth in them! But psychology has its own issues and so I want to try to contrast the two and then explore what the connections might be and how to resolve this kind of conflict between behavioral finance and traditional finance.

From my perspective—which admittedly is the perspective of a dilettante of psychology at this point—I observe that psychologists tend to base their analysis on observation and experimentation. In other words, there’s rarely a “theoretical” psychologist. Psychology, by its nature, starts with observation and then begins to generalize on the basis of those observations. Empirical analysis gives rise to theories, often resulting in multiple theories of behavior. So there isn’t just one model of the human mind or one model of a particular psychological phenomenon—there are multiple models. What’s interesting is that these models don’t have to be mutually consistent; they might be contradictory!

And that's somewhat puzzling to an economist because economists tend to prize internal consistency very highly.

And in a way, I think this makes sense, because, you know, as the old Almond Joy/Mounds commercial says, "sometimes you feel like a nut, and sometimes you don't." Humans are generally not consistent, so why should theories about their behavior be consistent? That's psychology. Now what about economics?

Economics, in contrast, is based almost entirely on abstraction and idealizations. You know the old joke about "assume a can opener?" Economists generally base all of their work on simplifying assumptions, regardless of whether or not those assumptions have anything to do with reality. Theories come first in economics, and accolades are routinely given to those economists with the most beautiful models. Whether or not such models have any direct bearing on reality often seems secondary.

I do think that empirical research is becoming more fashionable in economics, but for a long time, theorists occupied the forefront of our discipline. For example, field experiments are gaining more acceptance in the economics mainstream, but there is still considerable controversy as to whether economics is truly an experimental science. I certainly think it is, but I know many fellow economists who would strongly disagree!

Another important difference between psychology and economics is that there are relatively few theories of economic behavior. In fact, I really know of only one: maximize expected utility subject to a budget constraint. This is something that you're taught on the first day of basic microeconomics, and it's the *only* theory of economic behavior that I was ever exposed to in graduate school.

And finally, mutual and internal consistency of economic models is absolutely critical. If a model violates the implications of the standard canon, then it's simply not considered a good model by most economists. For example, if an economic model implies non-utility-maximizing individuals, then it will be almost impossible to get it published in a mainstream economics or finance journal. Maybe now it's less difficult, but certainly a few years ago it would have been unheard of.

Despite these differences, or perhaps because of them, I would like to argue that finance has revolutionized economics. This is one of most exciting aspects of finance in my view—at the end of the day, there's a bottom line. Someone once joked that if you take all the economists in the world and lay them end-to-end, they'd never reach a conclusion. But finance does, in fact, reach a conclusion because there are financial markets out there that impose a certain practical discipline to what we study and model in financial economics.

So this is my perspective on the contrast between economics and psychology. I hope you see the tensions

here, and the significant differences in the way that these two disciplines approach the same phenomena. I think there's a reason for these differences. Let me take a few minutes to discuss the *sociology* of economics and explain how it is that we got to this point, and then I'll turn to the main point of my talk—a new paradigm that reconciles behavioral and traditional finance. The reason that economics has gotten to its current state is because we economists, myself included, suffer from a peculiar psychological disorder known as "physics envy." I think you all know exactly what I'm talking about. We would love to have three laws that explain 99 percent of economic behavior; instead, we have about 99 laws that explain maybe 3 percent of economic behavior! Nevertheless, we like to talk as if we're dealing with physical phenomenon. And I have to admit, it gives me great pleasure to know that Sir Isaac Newton, the founding father of modern physics, lost 20,000 pounds in his investment in the South Sea Company, a classic stock market bubble. I take great comfort in the fact that this renowned physicist also succumbed to the temptations of foolish investments!

Why do we have physics envy? We can trace this phenomenon back to the source of most of modern economics: Paul A. Samuelson. Samuelson has done virtually everything notable in economics, including finance. In fact, he's one of the founders of what we now consider to be modern finance. In Samuelson's 1947 Ph.D. thesis, which he boldly called *Foundations of Economic Analysis*, he openly stated that his intention was to apply the principles of statistical thermodynamics to economics. In particular, to predict the path of a particle in a gravitational field, physicists came up with a concept called "action" and showed that the path followed by that particle in a gravitational field has the remarkable property that it minimizes the action. In other words, you can predict the trajectory of this particle by assuming that it moves through the field so as to minimize action. Of course, the particle is an inanimate object and doesn't really "know" how to minimize anything, but by modeling its behavior as if it does, physicists are able to predict the motion of such particles extremely precisely.

Well, Samuelson took a similar approach—he assumed that individuals acted so as to maximize a quantity called "expected utility" (actually, at the time he published his thesis, he focused only on the certainty case so he used "utility" and later extended his work to the uncertainty case). By modeling economic agents in this way, he hoped to be able predict their behavior in much the same way that physicists predicted the behavior of physical objects. From that point on, virtually everything that we have done in neoclassical economics has followed suit. And there's no doubt that we have made enormous progress thanks to Samuelson—as Newton once said "If I have seen further ... it is by standing upon the shoulders of Giants."

But what I have finally concluded, after years of frustration in attempting to reconcile the standard economics paradigm with the realities of financial markets, is that physics may not be the right metaphor for economic systems. In other words, the economy is not a physical system. This should be obvious to anyone who thinks about it for a moment—the focus of economic analysis is human interactions, not interactions of particles suspended in a fluid. And so when you're dealing with humans, physics becomes much less relevant. Another paradigm is needed.

Moreover, consider what drives the stock market. Economic fundamentals, to be sure. But not just those fundamentals, but also what I *think* of those fundamentals. And also what you think of those fundamentals. But also what I think you think of those fundamentals, and what you think I think you think of those fundamentals, and so on, ad infinitum, and ad nauseum! Such kinds of complex interactions really don't exist in physics. Physical laws don't change depending on what day it is or what happened in yesterday's markets. But economic laws do, at least the heuristics that we pretend to be economic laws.

Biology and Finance

My new paradigm—which I haven't completely formalized yet, but the various bits and pieces are now finally starting to fall in place—is that *biology* is a much better analogy to what we do in economics than physics, because biology faces the same kind of challenges that we face. Biologists deal with objects that don't always behave in ways that they can fully understand, and the kind of laws that biologists develop—with the exception of molecular biology, which I'll come back to in a few minutes—is probably of the same order of magnitude in explanatory power as some of the laws that economists have developed.

What I'm proposing is that we breathe some new life into an old paradigm, an old paradigm I'll turn to shortly which was developed by a well-known economist who died recently. I'm referring to Herbert Simon, who coined the term “bounded rationality” and started this literature several decades ago, and who received a Nobel Prize for this work. I'll argue that by looking at biological systems—especially from the perspective of evolutionary biology—we can actually recast much of the phenomena that we don't yet fully understand in a context that actually makes them much more transparent and logical.

That's a big claim. I'm not at all sure that I can defend it, which is why I'm giving a lunch talk as opposed to an hour seminar here, but let me give it a try.

At the same time, I want to emphasize that while we can criticize traditional finance and neoclassical economics; there are also issues with behavioral econom-

ics. The first is that behavioralists don't often talk about *why* behavioral biases exist. And I think that the “why” is probably more important than the “how,” partly because we've already spent a lot of time; many of you in the audience, in particular; documenting the “how.” Now that we've seen so many examples of these biases, it's time to ask the larger question: why are they there? I think I have an answer to that, but I leave you to judge whether or not it makes sense.

Second, how do these biases interact with the kind of institutional restrictions that I discussed earlier, restrictions like the Dutch book theorem? It's useful to have documented that individuals engage in “mental accounting,” but is there a reason for this and does the mental accounting really get in the way of decision-making from the perspective of a chief financial officer or a portfolio manager?

Third, how do biases interact with other competing biases? We know that there are several types of biases and it's fine to acknowledge that we may not have a single internally consistent model for all these biases, but at least we can acknowledge that they exist simultaneously and begin to explore their interactions. I'll give you an example of how to do this shortly.

And finally, what are their dynamics? How do behavioral biases change over time and with changing market conditions to create the kind of patterns like bubbles that we've documented in the empirical literature?

Emotion, Rationality and Heuristics

I think we're at the stage now where we can start answering these questions from a somewhat more formal context, and this new paradigm I'm proposing may be the first step in this direction. Towards that end, I'd like to refer to some recent research in the cognitive sciences, which is a little scary for me to be doing with Mike Gazzaniga and other real neuroscientists in the audience, but I'll remind you again, this is only a lunch talk!

I want to point out that we don't yet fully understand rationality to begin with. While there seems to be an obvious contrast between rationality and behavioralism, some recent evidence uncovered by cognitive neuroscientists seems to suggest that rationality and behavioral regularities are opposite sides of the same coin. This evidence is neatly laid out in a fascinating book titled *Descartes' Error* by Antonio Damasio, a neurologist at the University of Iowa. In this book, Damasio describes a patient of his who underwent surgery to remove a brain tumor. Along with the tumor, part of his frontal lobe had to be removed as well, and after the surgery, the big question on his and everybody else's mind was what kind of damage was done, and what kind of changes in his personality might result from the surgery.

They put him through a battery of diagnostic tests: IQ tests, language tests, mathematics tests, memory tests, the Minnesota multiphasic personality inventory, and several other tests, and they found nothing wrong with him. He was not impaired visually; his memory was just fine; his intelligence was average or above average; he could do arithmetic, logical inference, and all his mental faculties seemed pretty much intact. Yet within a few months of his surgery, he was fired from his job, his marriage fell apart, and his life was a mess. What had changed?

His doctors finally discovered one significant change in his personality—he was no longer capable of any emotional response. For example, although he remembered that prior to his surgery, he enjoyed a certain kind of music and certain types of food, after the surgery, he noted that he was unable to derive any pleasure from these things. He *knew* that these were things he had once enjoyed; yet now, he could *feel* nothing. This lack of emotional response had a pronounced affect on his day-to-day activities. Specifically, Damasio (1994, p. 36) writes:

When the job called for interrupting an activity and turning to another, he might persist nonetheless, seemingly losing sight of his main goal. Or he might interrupt the activity he had engaged, to turn to something he found more captivating at that particular moment. ... The flow of work was stopped. One might say that the particular step of the task at which Elliot balked was actually being carried out too well, and at the expense of the overall purpose. One might say that Elliot had become irrational concerning the larger frame of behavior ...

Apparently, Elliot's inability to feel, his lack of emotional response, made him act irrationally in his daily activities. This suggests that emotion plays a central role in rational behavior, which I find fascinating and completely counterintuitive. But upon further reflection, and in thinking back to Simon's paradigm of bounded rationality, this makes a great deal of sense. Emotion is what tells us when to stop focusing on the details and, instead, to turn to the larger task at hand. In doing so, emotion is apparently one of the primary drivers of rational behavior.

To illustrate this notion in more concrete terms, let me give you an example that involves something all of us do each day: getting dressed in the morning. Each morning, I have to select a specific set of clothing from my wardrobe, which consists of five jackets, 10 pairs of pants, 20 ties, 10 shirts, 10 pairs of socks, four pairs of shoes, and five belts. How many unique combinations do I have to choose from? Exactly 2,000,000 possible combinations! You might be jealous of that, but I guarantee you, you have at least as many in your own wardrobe! Now suppose it takes one second to evaluate each

of these combinations. In that case, a complete evaluation of all of my choices—which is required if I'm going to maximize the expected utility of my attire—would take 23.1 days! I can assure you that I actually get dressed a lot faster than that.

How are we able to get dressed so quickly each day? Obviously, it takes us only a few minutes to get dressed because we're not actually solving an optimization problem—we are not evaluating 2,000,000 combinations, but instead, we are engaged in what Simon calls "satisficing," heuristics that yield *satisfactory* solutions which are not necessarily optimal ones. In other words, we develop behavioral rules to solve this problem to a degree of optimality that is commensurate with the relative importance of the problem in the context of our daily lives. In contrast, solving the problem of getting dressed by computing the globally optimal outfit is supremely *irrational*. Yet this is exactly what Damasio's patient Elliot would attempt to do; after his surgery, he lost the capacity to *satisfice*.

Because there are many more important things in life than getting dressed, we solve this problem quickly and move on to the next task. We do this by developing heuristics and mental models that simplify these kinds of problems, and we adapt these heuristics and mental models from our experiences. In other words, we learn. This ability to learn, to adapt heuristics to changes in our environment and new information, is probably the single-most important evolutionary event to have occurred in the history of *Homo sapiens*. In fact, the very existence of what we call "civilization" is the direct outcome of this extraordinary adaptation, which is so remarkable that it has brought us to point where we now have at our disposal the means to change the very course of evolution itself. I'm referring to the fact that just this year, scientists announced the completion of the sequencing of the human genome. With that information and the technology that comes with it, we can determine the correspondence between genes and physiology or, eventually, behavior, and through the many breakthroughs in genetic engineering, we can change the genetic structure of organisms directly. This is a profound achievement—wonderful and terrible at the same time—and it is a direct result of our ability to learn and to pass on that learning from one generation to the next.

However, there are also costs associated with learning ability, and one of the most significant costs are behavioral biases that yield less than optimal outcomes when the learned behavior is taken out of context. One example is probability matching, which Mike Gazzaniga mentioned last night. Probability matching seems like a rather bizarre behavioral bias, especially for such intelligent animals as humans. After all, rats don't seem to suffer from such biases—why should *we*? The reason is that rats don't make forecasts; we do.

Let me give you a more concrete example. Yesterday, I was in the men's room on this floor, and while I was in there, somebody else walked in and tripped as he entered. I thought to myself, "that person is rather clumsy." But a minute later, someone else walked in and also tripped as he entered. At that point, I thought to myself "there must be something wrong with the flooring at the entrance." And sure enough, when I inspected the entrance, I noticed that the floor of the men's room was half a step lower than the entryway, and on the wall facing the doorway, there was a printed sign that read "Mind Your Step." I then concluded that this entrance was rather poorly designed.

Now think about the cognitive process that this experience represents. My first conjecture was that the first person to enter was clumsy; with only one additional data point, I changed my assessment of the situation entirely, rejecting the notion of "clumsy person" in favor of "poorly designed entrance." It turns out that this updated mental model was the correct one. But from a purely statistical point of view, to have one data point change your inference so dramatically seems highly suspect. In the right context, it's the smart thing to do. However, in the wrong context, it becomes a "behavioral bias," like mental accounting, loss aversion, overconfidence, and many of the other departures from "market rationality" that critics of traditional finance and economics like to cite.

Now let me return to the example of probability matching, a well-known behavioral bias in humans that has been documented in countless situations. One of the most interesting cases I've run across is the behavior of World War II bomber pilots who, as they boarded their bombers at the start of a mission, were allowed to wear either a flak jacket or a parachute, but because of the weight and bulkiness of both, pilots were allowed to choose only one. A flak jacket would protect the pilot from shrapnel in case the bomber was strafed by an enemy fighter's guns, and a parachute would be needed in case the bomber was shot down. Now as I recall, there was a 75 percent chance of getting strafed and a 25 percent chance of getting shot down, so the rational choice would have been for the pilots to select the flak jacket. However, military records show that bomber pilots tended to choose flak jackets 75 percent of the time and parachutes 25 percent of the time—in other words, they engaged in probability matching.

Now, obviously, over time, and with enough experience, individuals will eventually overcome these biases, as Mike Gazzaniga pointed out in his talk last night. But the question of the time scale is very critical here—given the natural time scale of the phenomenon at hand, will individuals be able to learn that there is no predictability and then change their behavior accordingly?

As a final thought regarding probability matching, let me share with you an interesting fact I came across

recently while browsing a neurosciences journal. I read an article concerning the frontal cortex of the chimpanzee, and purely as an aside, and buried in a footnote was a comment by the authors that they had to conduct certain training sessions with their chimpanzees prior to their main experiments because they discovered that without this additional training, their chimps engaged in probability matching, and this would have been problematic for their particular experimental design. Rats don't engage in probability matching, but chimpanzees do! Although this is by no means proof that probability matching is a concomitant of higher mental faculties and learning ability, it is certainly suggestive.

Satisficing: An Evolutionary Perspective

The ability to learn about and adapt to a changing environment is a critically important aspect of human evolution. It's a revolutionary adaptation, but it has its costs. In the context of financial markets, if you've been a value manager over the past five years, then you've been losing money five years in a row, and that has to make you wonder whether or not you are badly mistaken and barking up the wrong tree. There's nothing wrong with that, and you may be wrong to think that you're barking up the wrong tree, but the fact is that your ability to forecast, your ability to learn will cause you to question yourself constantly, even when you're right.

Now I would argue that Simon pointed this out years ago, perhaps not in so many words, and not with the same buttressing arguments from the cognitive sciences, but he did suggest that individuals don't optimize, they satisfice, they optimize until it's "close enough." But how close is "close enough?"

This was perhaps the most serious stumbling block for Simon's seminal contribution to economics, and as a result, there has been relatively little research devoted to satisficing since its introduction almost half a century ago. Specifically, if it's impossible to evaluate the 2,000,000 combinations of my wardrobe to determine the optimal outfit, how do I go about satisficing? To balance the costs and benefits of evaluating a certain number of possible solutions, don't we need to know what the optimal solution is? Otherwise, how do we know how close we are to the optimum and how to trade off further optimization against the cost of computation? And if we don't know, well, then, how do we satisfice?

And that's where the literature has stayed for over 40 years or so. It wasn't until last year that the answer finally dawned on me—maybe this is already out there in the literature, but I haven't seen it yet. I certainly hope it isn't, because then I can't claim to have discovered it! The answer is simply this: we *don't* balance the

costs and benefits. We evolve and adapt. Adaptation and evolution dictate the balance of bounded rationality, and it's the process of natural selection that determines where you are in the spectrum of how much you optimize.

This principle can be applied to many economic or financial contexts, including biases in earnings forecasts, the birth and death of financial bubbles, probability matching, overconfidence, overreaction, loss aversion, and other situations where behavior and rationality seem to be at odds with each other. Now, is this the right answer?

Well, I don't know. There is a great deal of research that needs to be done before we can conclude that this evolutionary approach is the right one, but it does close the system. It provides a satisfying resolution—at least for me—of the near-religious debate between Simon and the rest of the economics profession regarding the possibility of satisficing without first optimizing. The answer is simply trial-and-error and natural selection. And the institutional and market forces that I referred to earlier, those are the forces that create competitive pressure and provide the natural selection, which determines the degree of optimization that is accomplished.

The Power of Evolution

Now the last topic I would like to turn to is an illustration of the forces of evolution, to give you some idea of how powerful this new paradigm can be in explaining change and innovation. It's a very simple example, which has nothing to do with satisficing, but I think you'll see its relevance to our current discussion immediately. I'd like to ask you to think about an experiment in which we randomly select nine letters of the English alphabet. What do think the chances are that such a draw will produce the phrase "CAPMLIVES?" Well, the odds are one in $26^9 = 5,429,503,678,976$, approximately one in five-and-a-half trillion; very unlikely indeed.

But I'm going to propose a little evolution in this example. Let's start with the random draw "TJIXOMSFZ" and suppose that this particular draw has 10 "offspring." By this I mean that 10 new draws are generated from "TJIXOMSFZ," and here are the rules by which each offspring is produced: with 99 percent probability, each letter of an offspring will be the same as its parent, and with 1 percent probability it will change or "mutate." Therefore, the children will look very much like their parent. If a letter does mutate, it will change to any of the other 25 letters with equal probability. So to summarize, there's a one out of 100 chance that a mutation will occur for each letter, and if it does occur, there is a 1 in 25 chance that the mutation will be "beneficial" (in other words, the mutation

will create a match to the corresponding letter in "CAPMLIVES"). And let's assume that these random draws are independently and identically distributed for each letter of an offspring, and for all offspring. Finally, impose the following selection process for "survival of the fittest": Of the 10 offspring produced in each generation, the only one that survives to "reproduce" is the offspring with the largest number of letters matching "CAPMLIVES" (in case of ties, we'll simply randomize). Then the process begins again, with the surviving offspring producing 10 children of its own under the same mechanism and with the same probabilities.

In this simple set-up, there are many interesting evolutionary quantities that can be computed. For example, what's the probability that at least one offspring has at least one letter matching CAPMLIVES? It turns out that the probability of at least one match is about 3.5 percent. How about the average number of generations required for at least one match occurs? About 28 generations. But the main question I'd like to put forward is this: How many generations do you think it would require for this type of evolutionary process to yield "CAPMLIVES?" A thousand? A million? This seems like a very unlikely event indeed. However, the answer is about 751 generations on average, a remarkably small number of generations compared with the one in five-and-a-half trillion chance of drawing "CAPMLIVES" randomly, especially in light of the trial-and-error nature of our evolutionary process. But this is the point of the example: evolutionary forces, despite their inefficiency, work inexorably towards certain objectives, and when those objectives are clearly defined by specific performance criteria such as "CAPMLIVES," optimal solutions may arise surprisingly quickly.

How long has *Homo sapiens* been around? We've been in existence for about a hundred thousand years. And if you assume that a typical generation is about 20 years, we've had approximately 5,000 generations to improve ourselves. This means that we've had many opportunities to develop satisficing algorithms to deal with an ever-changing environment, and that these heuristics are quite valuable in their proper contexts. But when taken out of context, they can be ineffective or, worse, detrimental.

Conclusions

So where does all this leave us? Let me conclude by telling you a bit about my current research agenda, which takes as its starting point the issues I've raised in this lunch talk. First, I would like to derive the evolutionary underpinnings of behavioral finance, extending Simon's notion of bounded rationality by developing the mechanism through which natural selection gives

rise to satisficing heuristics. Second, I would like to develop a deeper understanding of the statistical properties of selection. Several years ago, I wrote an article on the effects of data-snooping on statistical tests of financial asset-pricing models, and concluded that even small selection biases can have an enormous impact on test statistics and the inferences that are drawn from those statistics. Well, the flip side of that coin is that very small evolutionary biases can have a huge impact on the course of evolution and the satisficing heuristics that emerge. The last set of topics that I'm pursuing lie more in the domain of cognitive neurosciences proper, and these have to do with the psychophysiology and, ultimately, brain functions involved in real-time financial risk processing. To tackle these issues, I plan to study the psychology and physiology of financial traders, both in the laboratory and during live trading sessions.

But since I'm now out of time, let me finish by pointing out that the Harvard evolutionary biologist, Edward O. Wilson, founded a discipline known as "sociobiology" over 30 years ago—now known as "evolutionary psychology"—in which he applied the principles of evolution, not at the genetic level, but with respect to behavior and culture. Therefore, the ideas that Simon espoused in the 1950's have firm roots in evolutionary biology. In my view, the evolutionary perspective may be a more promising way to view human decision making, and I would encourage all of you to take a look at Wilson's recent book *Consilience*, in which he argues that over time, there will be a natural "jumping together" of various disparate disciplines. I think that's what we're witnessing today, disciplines such as psychology, the cognitive neurosciences, economics, and finance all coming together at certain common points of contact.

In Wilson's autobiography, *Naturalist*, he describes the bitter controversy and debate between evolutionary and molecular biologists at Harvard during the 1960's, as well as the happy rapprochement that the two fields have enjoyed recently. My hope is that 10 or 20 years from now, traditional and behavioral finance will reach the same satisfying resolution.

General Discussion

Voice: I'm concerned about the distinction between explaining and predicting, and I'm thinking that adaptation is always about adapting to the truth, which is in the past, but we're trying to know the truth for the future. I mean is this going to be a better explanatory tool or is it actually going to help us to forecast some day?

Andrew Lo: I believe that every mutual fund prospectus contains the phrase "past returns do not guarantee future performance," and I always found

that odd because the main reason that people want to know about past performance is to infer something about future returns! But you raise a good point. What I'm focusing on is not better forecasting methods. Frankly, I think we have pretty good forecasting methods as it is. These methods are unlikely to improve without adding more structural information about the market. What I'm trying to discover in my own research is this underlying structure, and one of my points is that this underlying structure is not "value," it's not "growth," and it's not the CAPM. The underlying structure is human behavior and how behavior interacts with institutions. And if you can develop a better understanding of this underlying structure, then I think you'll have a better chance of forecasting when "value" will be in favor and when "growth" will be in favor.

Aaron Lynch: My understanding is that Wilson emphasized molecules in a way as being the locus of the information that's being selected. In other words, the variants that are being selected are DNA, and overlooks to a large extent the effect of ideas, and variants of ideas. For example let's consider fashion. A number of ways to dress come out and one gets a little more attention and more social status. As this process is reiterated the locus of information is neurally stored instead of molecularly stored.

Andrew Lo: Yes, this is exactly the point that the evolutionary psychology literature is making.

Voice: But my understanding of that is there's a tendency to emphasize genes.

Andrew Lo: To some degree, all of these issues come down to genes, right?

Voice: Well, genes would be underlying the process, but they wouldn't be in the actual selection of the fashion, for example. There are no genes for some cultural issues.

Andrew Lo: I think the answer is that there's selection going on at many different levels.

Voice: And the idea of wearing a jacket or a tie?

Andrew Lo: Cultural norms. One of my colleagues at MIT, Steven Pinker, has just completed a rather controversial paper—I don't know if it's published yet—on religion, in which he asks whether a society without any notion of God would eventually create such a notion to satisfy certain basic needs that humans have as part of their cognitive development? He makes a fascinating argument that religion is an outcome of evolutionary forces.

Voice: You see, I make the argument that if you have a religion that's polytheistic; it will evolve before its monotheism.

Voice: You know, since we know where most of our evolution took place, we can go to the next step

and simply look at how they would be optimized and then observe the malfunctions that occur when they play the stock market.

Andrew Lo: Well, that's one of the things that I've actually been working on, although not exactly in those terms. Let me explain. There's a very large literature known as "behavioral ecology" in which optimization models are used to develop implications for behavior. For example, the herding behavior of springbok grazing on an Africa veldt can be shown to maximize the probability of survival given the way that lions tend to hunt these springbok. Along the same lines, I've been able to derive the familiar S-shaped utility function that Kahneman and Tversky document as "loss aversion" by appealing to optimization principles in an evolutionary setting. Specifically, S-shaped Utility functions emerge naturally as ones that maximize the probability of survival under certain types of selection pressures. The intuition is clear: an S-shaped curve implies that if you're threatened with extinction, you'll be extremely risk-seeking and anything goes, but once you reach the level of safety, you become extraordinarily conservative. Someone once said that if you're young and conservative, you have no heart, and if you're old and liberal, you have no brain! Apparently, there may be some evolutionary basis for such a statement!

Arnold Wood: Andrew, thank you very much.

Acknowledgments

This is an edited version of a luncheon address delivered on June 1, 2001 at the Third Annual Institute of Psychology and Markets Conference. I thank Gunduz Caginalp, David Dreman, Arnie Wood and conference participants at the Third Annual Institute on Psychology and Financial Markets Conference for helpful comments and discussion. Research support from the MIT Laboratory for Financial Engineering is gratefully acknowledged.

References

- Breeden, D. "An Intertemporal Capital Pricing Model with Stochastic Investment Opportunities," *Journal of Financial Economics*, 7, (1979), pp. 265–296.
- Damasio, A. *Descartes' Error: Emotion, Reason, and the Human Brain*. New York: Avon Books, 1994.
- Lucas, R. "Asset Prices in an Exchange Economy," *Econometrica*, 46, (1978), pp. 1429–1446.
- Merton, R. "An Intertemporal Capital Asset Pricing Model," *Econometrica*, 41, (1973), pp. 867–887.
- Samuelson, P. *Foundations of Economics Analysis*. Cambridge, MA: Harvard University Press, 1947.
- Simon, H. *Models of Bounded Rationality Volumes 1 and 2*. Cambridge, MA: MIT Press, 1982.
- Wilson, E. *Sociobiology: The New Synthesis*. Cambridge, MA: Belknap Press of Harvard University Press, 1975.
- Wilson, E. *Naturalist*. Washington, DC: Island Press, 1994.
- Wilson, E. *Consilience: The Unity of Knowledge*. New York: Alfred A. Knopf, 1998.

